

# From sand to circuits

How Intel makes integrated circuit chips

Sand with Intel® Core™2 Duo processor.



The Intel® Core™2 Duo processor with Intel 45nm High-k metal gate silicon technology

"The implementation of High-k and metal materials marks the biggest change in transistor technology since the introduction of polysilicon gate MOS transistors in the late 1960s."

Gordon Moore, Intel Co-Founder





# Revolutionary

They are small, about the size of a fingernail. Yet tiny silicon chips such as the Intel® Core™2 Duo processor that you see here are changing the way people live, work, and play.

The task of making chips like these—the most complex devices ever manufactured—is no small feat. A sophisticated chip, such as a *microprocessor*, can contain hundreds of millions or even billions of *transistors* interconnected by fine wires made of copper. These transistors act as switches, either preventing or allowing electrical current to pass through. A positive charge fed to a transistor's *gate* attracts electrons. This gate creates a channel between the transistor's *source* and *drain* through which electrical current flows, creating an “on” state. A negative charge at the gate prevents the current from being able to flow through, creating an “off” state for the transistor.

Intel uses its advanced manufacturing technology to build several hundred trillion transistors every day. Intel's breakthrough *45-nanometer (nm) High-k silicon technology* enables the production of transistors that are so small that 2 million of them would fit into the period at the end of this sentence.

Intel's success at reducing transistor size and maximizing performance results in advanced processor technology that helps drive other innovations in almost all industries. Today, silicon *chips* are everywhere—powering the Internet, enabling a revolution in mobile computing, automating factories, enhancing cell phones, and enriching home entertainment. Silicon is at the heart of an ever expanding, increasingly connected digital world.

Explore this brochure to learn how Intel makes silicon chips. If you are unfamiliar with a technical term, see the “Terminology” section at the end of the brochure. It defines the words that are *italicized* in the text.



# Design

Silicon chip manufacturing starts with a design, or a blueprint. Intel considers many factors. What type of chip is needed and why? How many transistors can be built on the chip?


What is the optimal chip size? What technology will be available to create the chip? When does the chip need to be ready? Where will it be manufactured and tested?

To answer these questions, Intel teams work with customers, software companies, and Intel's marketing, manufacturing, and testing staff. Intel design teams take this input and begin the monumental task of defining a chip's features and design.

When the specifications for the chip are ready, Intel creates a logic design, an abstract representation of the millions of transistors and interconnections that control the flow of electricity through a chip. After this phase is complete, designers create physical representations of each layer of the chip and its transistors. They then create stencil-like patterns, or *masks*, for each layer of the chip. Masks are used with ultraviolet light during a fabrication process called *photolithography*.

To complete the design, testing, and simulation of a chip, Intel uses sophisticated *computer-aided design (CAD)* tools. CAD helps designers create very complex designs that meet functional and performance goals.


After extensive modeling, simulation, and verification, the chip is ready for fabrication. It can take hundreds of engineers working full time for more than two years to design, test, and ready a new chip design for fabrication.



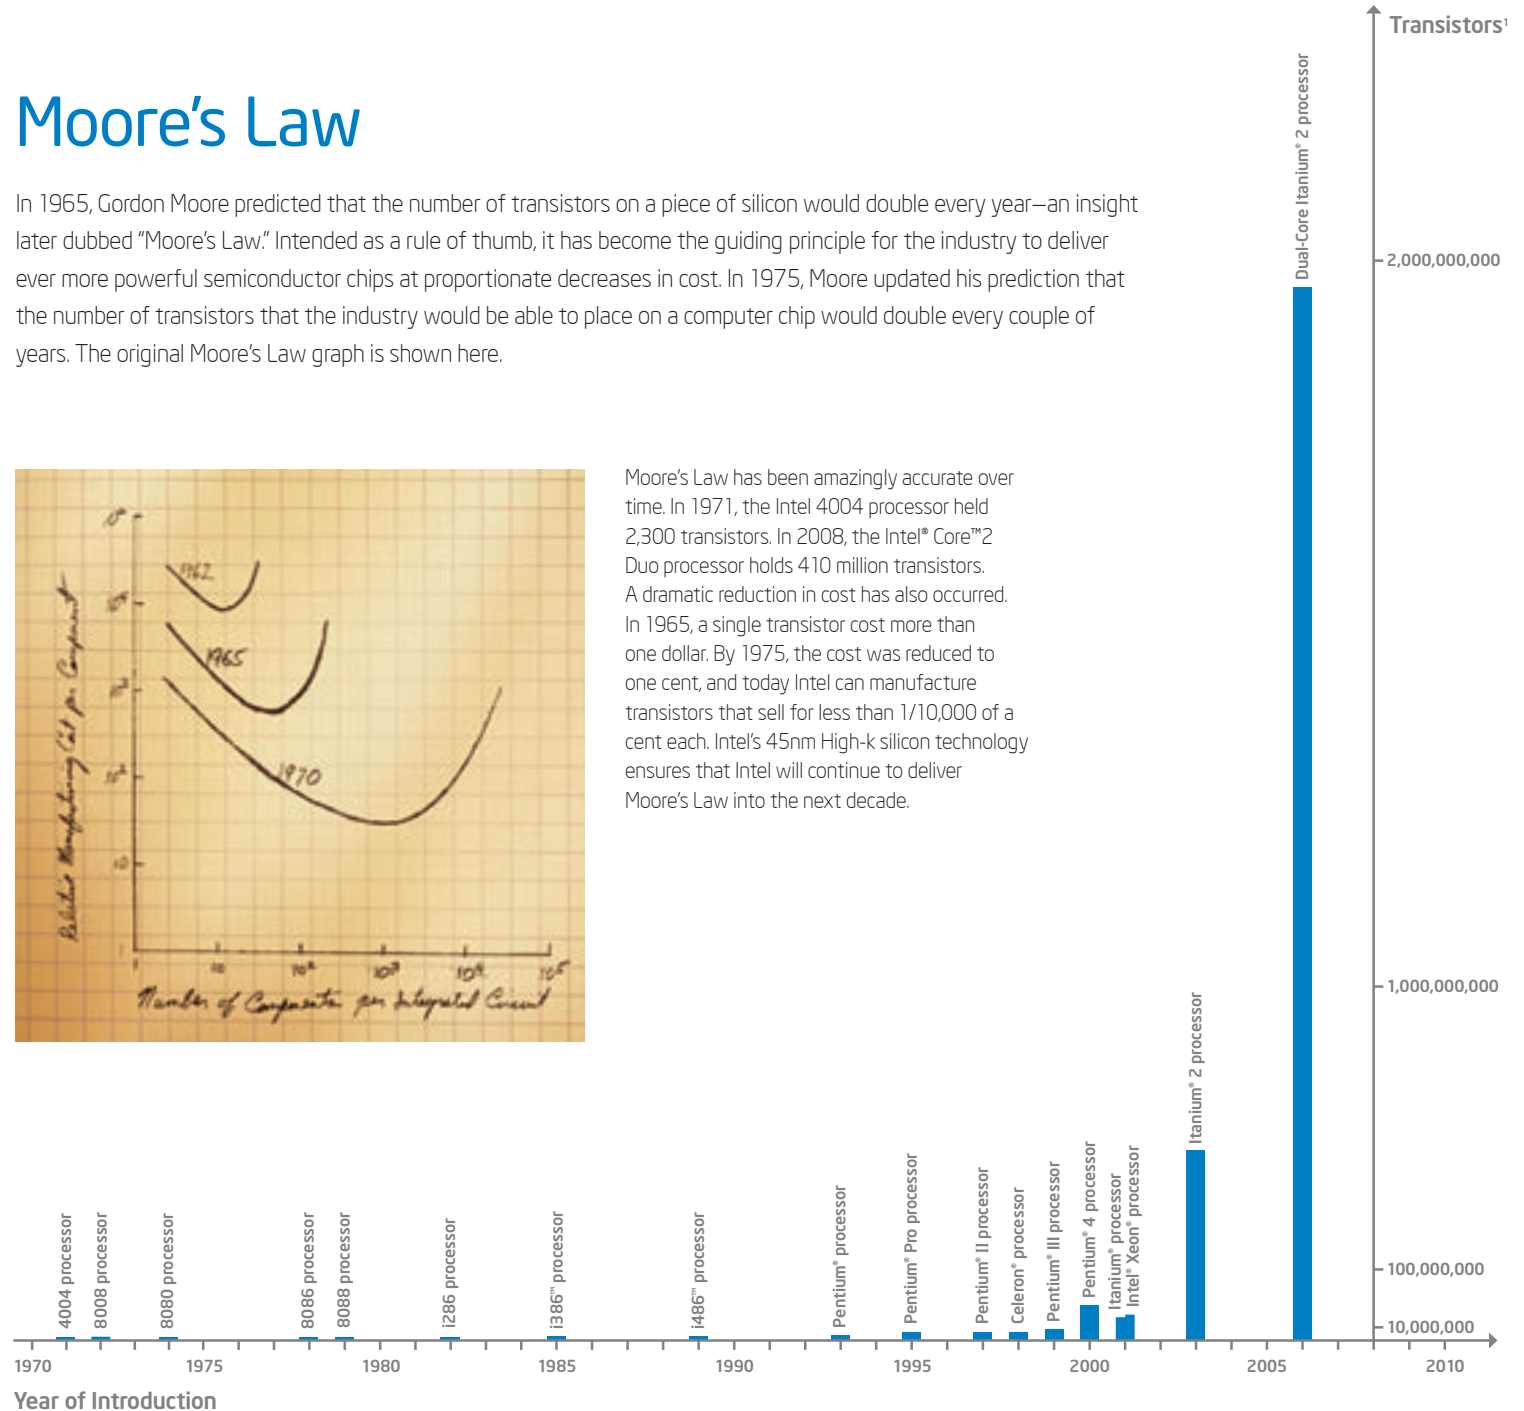
Intel creates a logic description of a new chip. Simulators take the register transfer level (RTL) code (shown in the background of this image) to verify the accuracy of the specification before fabrication begins.

# Moore's Law

In 1965, Gordon Moore predicted that the number of transistors on a piece of silicon would double every year—an insight later dubbed “Moore’s Law.” Intended as a rule of thumb, it has become the guiding principle for the industry to deliver ever more powerful semiconductor chips at proportionate decreases in cost. In 1975, Moore updated his prediction that the number of transistors that the industry would be able to place on a computer chip would double every couple of years. The original Moore’s Law graph is shown here.



Moore’s Law has been amazingly accurate over time. In 1971, the Intel 4004 processor held 2,300 transistors. In 2008, the Intel® Core™2 Duo processor holds 410 million transistors. A dramatic reduction in cost has also occurred. In 1965, a single transistor cost more than one dollar. By 1975, the cost was reduced to one cent, and today Intel can manufacture transistors that sell for less than 1/10,000 of a cent each. Intel’s 45nm High-k silicon technology ensures that Intel will continue to deliver Moore’s Law into the next decade.



<sup>1</sup> Vertical scale of illustration not proportional to actual transistor count.




# Fabrication

The process of making chips is called *fabrication*. The factories where chips are made are called fabrication facilities, or *fabs*. Intel fabs are among the most technically advanced manufacturing facilities in the world. Within these sophisticated fabs, Intel makes chips in a special area called a *cleanroom*.

Because particles of dust can ruin the complex circuitry on a chip, cleanroom air must be ultra-clean. Purified air is constantly re-circulated, entering through the ceiling and exiting through floor tiles.

Technicians put on a special suit, commonly called a bunny suit, before they enter a cleanroom. This helps keep contaminants such as lint and hair off the wafers. In a cleanroom, a cubic foot of air contains less than one particle measuring about 0.5 micron (millionth of a meter) across. That's thousands of times cleaner than a hospital operating room.

Automation also plays a critical role in a fab. Batches of wafers are kept clean and processed quickly and efficiently by traveling through the fab inside *front-opening unified pods (FOUPs)* on an overhead monorail. Each FOUP receives a barcode tag that identifies the recipe that will be used to make the chips inside. This labeling ensures the correct processing at each step of fabrication. Each FOUP contains up to 25 wafers and weighs more than 25 pounds. Production automation machinery allows for this FOUP weight, which is too heavy to be handled manually by technicians.



1




2



3

- 1 Orange FOUPs carry 300mm wafers in an automated fab.
- 2 Highly trained technicians monitor each phase of chip fabrication.
- 3 Purified air enters from the ceiling and exits through perforated floor tiles.
- 4 A technician holds a 300mm wafer.



The “recipe” for fabricating a chip varies depending on the chip’s proposed use. Intel uses a variety of ingredients and performs as many as 300 steps with chemicals, gas, or light to complete fabrication.

#### A sandy start


It all starts with *silicon*, the principal ingredient in common beach sand. Intel builds chips in batches on *wafers* made of ultra-pure silicon. Silicon is a *semiconductor*. This means that unlike insulators such as glass (which always resist the passage of electrons) or conductors such as copper (which generally let electrons pass through), silicon can be altered to be a conductor or an insulator. Silicon is a good choice for making wafers because it is abundant, its oxide is a good insulator, and the industry has decades of experience working with it.

To make wafers, silicon is chemically processed so that it becomes 99.9999% pure. The purified silicon is melted and grown into long, cylindrical *ingots*. The ingots are then sliced into thin wafers that are polished until they have flawless, mirror-smooth surfaces. When Intel first started making chips, the company used 2-inch-diameter wafers. Now the company uses primarily 12-inch, or 300-millimeter (mm) wafers; larger wafers are more difficult to process, but the result is lower cost per chip.

#### Layer by layer


Intel uses a photolithographic “printing” process to build a chip layer by layer. Many layers are deposited across the wafer and then removed in small areas to create transistors and interconnects. Together, they will form the active (“on/off”) part of the chip’s circuitry plus the connections between them, in a three-dimensional structure. The process is performed dozens of times on each wafer, with hundreds or thousands of chips placed grid-like on a wafer and processed simultaneously.

**Silicon, the principal ingredient in beach sand, is a natural semiconductor and the most common element on earth after oxygen.**




# Process


Building *circuits* to form a computer chip is extremely precise and complex. It requires dozens of layers of various materials in specific patterns to simultaneously produce hundreds or thousands of *die* on each 300mm wafer. The following illustration takes a closer look at the process of adding one layer—a single patterned oxide film.




1. Start with a partially processed die on a silicon wafer.



2. Deposit oxide layer.



3. Coat with photoresist.



4. Position mask and flash ultraviolet light.


- 1. Start with a partially processed die on a silicon wafer.** A chip is often referred to as die until final packaging has been completed.
- 2. Deposit oxide layer.** A thin film of *oxide* is an electrical insulator. Like the insulator surrounding household wires, it is a key component of electronic circuits. Intel “grows” this layer of oxide on top of the wafer in a furnace at very high temperatures in the presence of oxygen.
- 3. Coat with photoresist.** A light-sensitive substance called *photoresist* prepares the wafer for the removal of sections

of the oxide to create a specific oxide pattern. Photoresist is sensitive to ultraviolet light, yet it is also resistant to certain etching chemicals that will be applied later.


- 4. Position mask and flash ultraviolet light.** Masks—pieces of glass with transparent and opaque regions—are a result of the design phase and define the circuit pattern on each layer of a chip. A sophisticated machine called a stepper aligns the mask to the wafer. The stepper “steps” across the wafer, stopping briefly at incremental locations to flash ultraviolet light through the transparent regions of the mask. This process is called photolithography. The portions of the photoresist that are exposed to light become soluble.

- 5. Rinse with solvent.** A solvent removes the exposed portions of photoresist, revealing part of the oxide layer underneath.
- 6. Etch with acid.** Using an acid in a process called *etching*, the exposed oxide is removed. Oxide protected by the mask remains in place.
- 7. Remove remaining photoresist.** Finally, the remaining photoresist is removed, leaving the desired pattern of oxide on the silicon wafer. A new oxide layer is complete.







A completed die contains millions of circuits that appear as an intricate pattern.




5. Rinse with solvent.



6. Etch with acid.



7. Remove remaining photoresist.



Magnified cross-section of metal interconnects with oxide layers.

## Performing More Fabrication Steps

Laying down an oxide layer is just one part of the fabrication process. Other steps include the following.

### Adding more layers

Additional materials such as *polysilicon*, which conducts electricity, are deposited on the wafer through further film deposition, masking, and etching steps. Each layer of material has a unique pattern.

### Doping

The *doping* operation bombards the exposed areas of the silicon wafer with various chemical impurities, altering the way the silicon in these areas conducts electricity. Doping is what turns silicon into silicon transistors, enabling the switching between the two states, on and off, that represent *binary* 1s and 0s, which provide the basis for representing information in a computer.

### Metallization

Multiple layers of metal are applied to form the electrical connections between the transistors. Intel uses eight or

more patterned layers of copper because of its low resistance and because it can be cost-effectively integrated into the manufacturing process. Interconnects between layers, called contacts, are made of tungsten. The specific patterns of these metals are also formed using photolithography, as described previously.

### Completing the wafer

A completed wafer contains millions or even billions of transistors connected by a multi-layer maze of metal "wires." Finally, the wafer is coated with a *passivation* layer to help protect it from contamination and increase its electrical stability.



# Testing and Packaging

After creating layers on the wafers, Intel performs *wafer sort*, and a computer completes a series of tests to ensure that chip circuits meet specifications to perform as designed.

Intel sends the approved wafers to an Intel assembly facility, where a precision saw separates each wafer into individual rectangular chips, called die. Each functioning die is assembled into a package that, in addition to protecting the die, delivers critical power and electrical connections from the main circuit board on a computer. It is this final “package” that is placed directly on a computer circuit board or in other devices such as cell phones and personal digital assistants (PDAs).

As processor technologies advance, the demands on packaging to support and optimize the technologies increase. Because Intel makes chips that have many different applications, the company uses a variety of packaging technologies.

## High-performance packages

**Flip-chip packaging.** *Flip-chip packaging* is an example of one of the advanced packages that Intel uses. To package the die, Intel begins by attaching tiny metal bumps on the die surface to the supporting base, or substrate of the package, completing an electrical connection from the chip to the package. This method is called “flip chip” because the silicon die are “flipped” to their front side for attachment, compared to other types of packaging that attach to the back of the die. Intel uses an organic or polymer substrate to enable higher performance copper electrical interconnections from the die to the circuit board. A compliant material is then added between the substrate and the die to manage mechanical stress. In the last step, Intel attaches a structure called a heat spreader to help disperse the heat generated by the chip during normal use.



**Wire bond for stacked-chip packaging.** *Stacked-chip packaging* technologies result in packages that are only slightly larger than the multiple silicon die that they contain. Intel stacks multiple memory and logic die in a single package to increase performance and minimize the use of space, which are critical in today's small handheld devices. When attaching the die, Intel uses a special material that is optimized for mechanical, thermal, and electrical performance to “glue” the first die to the substrate. The other die are then stacked and “glued” to each other to create a combination of chips that meet product performance goals.

After the die are attached, sophisticated tools bond extremely fine wires from each die to the substrate. This process, called *wire bonding*, is repeated for each die included in the stack until all die are electrically connected to the same package. The die are then encapsulated with a molding process and a protective coating that flows into the narrow spaces between the die and the package. Lastly, Intel attaches specialized alloy “balls” to the bottom of the package to electrically connect the package to the circuit board.


## One more check

Intel performs reliability and electrical “tests” on each completed unit. The company verifies that the chips are functional and perform at their designed speed across a variety of temperatures. Because chips may end up in items ranging from automobile engines to spacecraft and laptops, they must be able to withstand many different environmental stresses. Chips are also tested for long-term reliability to ensure that they will continue to perform as specified. Upon approval, chips are electrically coded, visually inspected, and packaged in protective materials for shipment to Intel customers.






Front



Back



The tiny bumps on this package provide the electrical connections between the chip and the circuit board.

A technician inserts a circuit board containing many packages.







# Innovation

Intel's processor technologies offer exciting advancements, including state-of-the-art chips with multiple cores, or "brains."

These brains enable the efficient execution of parallel tasks, such as when a computer simultaneously performs word processing, plays music, prints a file, and checks for viruses. *Multi-core processor* architectures significantly improve performance while increasing energy efficiency, which is an important consideration in today's high-performance products.

Intel has a long history of translating technology leaps into tangible benefits. It's not just about making technology faster, smarter, and cheaper—it's about using that technology to make life better and our experiences richer.







Intel's advanced chips and manufacturing are helping to bring together the best of computing, communications, and consumer electronics to enable broader and ever more valuable benefits from technology.





# Undisputed Leadership

Unwavering commitment to moving technology forward

1970s



**4004 Processor**  
Introduced: 1971




**8008 Processor**  
Introduced: 1972



**8080 Processor**  
Introduced: 1974



**8086 Processor**  
Introduced: 1978



**8088 Processor**  
Introduced: 1979

The 4004 microprocessor delivered the same computing power as the first electronic computer, the ENIAC; which filled an entire room.

1980s



**Intel286 Processor**  
Introduced: 1982



**Intel386™ Processor**  
Introduced: 1985



**Intel486™ Processor**  
Introduced: 1989

Intel386™ microprocessors were built into engineering workstations, PCs, and network file servers.

1990s



**Pentium® Processor**  
Introduced: 1993




**Pentium® Pro Processor**  
Introduced: 1995



**Pentium® II Processor**  
Introduced: 1997



**Celeron® Processor**  
Introduced: 1998




**Pentium® III Processor**  
Introduced: 1999

The Pentium® processor brought faster performance, better graphics, and real-time speech and video to personal computers.

2000s



**Pentium® 4 Processor**  
Introduced: 2000




**Itanium® Processor**  
Introduced: 2001



**Intel® Xeon® Processor**  
Introduced: 2001



**Itanium® 2 Processor**  
Introduced: 2003



**Dual-Core Itanium® 2 Processor**  
Introduced: 2006



**Intel® Core™2 Duo Processor**  
Introduced: 2008

The 45nm High-k Intel® Core™2 Duo processor contains 410 million transistors. The Intel Core 2 Duo processor was first introduced in 2006 at 65nm, containing 291 million transistors.





# Terminology

**Intel 45nm High-k metal gate silicon technology:** One of the biggest advances in fundamental transistor design. Intel's innovative combination of metal gates and High-k gate dielectrics reduces electrical current leakage as transistors get ever smaller.

**Binary:** Having two parts. The binary number system that computers use is composed of the digits 0 and 1.

**Channel:** The region under the gate of a transistor where current flows when the transistor is in the "on" state.

**Chip:** A tiny, thin square or rectangle that contains integrated electronic circuitry. Die are built in batches on wafers of silicon. A chip is a packaged die. See also "Microprocessor."

**Circuit:** A network of transistors interconnected by wires in a specific configuration to perform a function.

**Cleanroom:** The ultra-clean room where chips are fabricated. The air in a cleanroom is thousands of times cleaner than that in a typical hospital operating room.

**Computer-aided design (CAD):** Sophisticated computerized workstations and software that Intel uses to design integrated circuits.

**Die:** Alternate name for a chip, usually before it is packaged. See also "Chip."

**Doping:** A wafer fabrication process in which exposed areas of silicon are bombarded with chemical impurities to alter the way the silicon in those regions conducts electricity.

**Drain:** A highly doped region adjacent to a transistor's current-carrying channel that transports electrons from the transistor to the next circuit element or conductor.

**Etching:** The removal of selected portions of materials to define patterned layers on chips.

**Fab:** A shortened term for "fabrication facility," where Intel manufactures silicon chips.

**Fabrication:** The process of making chips.

**Flip-chip packaging:** A type of chip package in which a die is "flipped" to its front side and attached to the package, compared to packaging such as wirebond that attaches the back of the die to the package.

**Front-opening unified pod (FOUP):** A container that holds and carries wafers as part of an automated system in a fab.

**Gate:** The input control region of a transistor where a negative or positive charge is applied.

**Gate dielectric:** A thin layer underneath the gate that isolates the gate from the channel.

**High-k material:** A material that can replace silicon dioxide as a gate dielectric. It has good insulating properties and creates a high field effect between the gate and channel. Both are desirable properties for high-performance transistors. Also, because High-k materials can be thicker than silicon dioxide, while retaining the same desirable properties, they greatly reduce current leakage.

**Mask:** A stencil-like pattern used during fabrication to "print" layered circuit patterns on a wafer.

**Microprocessor:** The "brain" of a computer. Multiple microprocessors working together are the "hearts" of servers, communications products, and other digital devices. See also "Chip."

**Multi-core processor:** A chip with two or more processing cores, or "brains."

**Nanometer:** One billionth of a meter.

**Oxide:** An insulating layer that is formed on a wafer during chip fabrication. Silicon dioxide is one example.

**Passivation:** The process of coating a silicon chip with an oxide layer to help protect it from contamination and increase its electrical stability.

**Photolithography:** The process of creating a specific pattern of material onto a silicon wafer by using ultraviolet light and a mask to define the desired pattern.

**Photoresist:** A substance that becomes soluble when exposed to ultraviolet light. Analogous to photographic film, it is sensitive to ultraviolet light but is also resistant to certain etching chemicals. Used to help define circuit patterns during chip fabrication.

**Polysilicon:** A shortened term for "polycrystalline silicon," or silicon made up of many crystals. This conductive material is used as an interconnect layer on a chip, and as the gates of transistors.

**Register Transfer Level (RTL) code:** A computer language that processor designers use to create a functional description of the chip. RTL is used to define, simulate, and test processor functionality, before actually producing the processor.

**Semiconductor:** A material (such as silicon) that can be altered to conduct electrical current or block its passage.

**Silicon:** The principal ingredient in common beach sand and the element used to make the wafers upon which chips are fabricated. It is a natural semiconductor and is the most common element on earth after oxygen.

**Silicon ingot:** A cylinder formed of 99.9999% pure silicon.

**Source:** The region of a transistor where electrons move into the channel.

**Stacked-chip packaging:** A type of chip package that contains multiple die stacked in a single package.

**Transistor:** A type of switch that controls the flow of electricity. A chip may contain millions or billions of transistors.

**Wafer:** A thin silicon disc sliced from a cylindrical crystal ingot. Used as the base material for building integrated circuits.

**Wafer sort:** An electrical test procedure that identifies the chips on a wafer that are not fully functional.

**Wire bonding:** The process of connecting extremely thin wires from a chip's bond pads to leads on a package.



[www.intel.com](http://www.intel.com)

Learn more about Intel history at [www.intel.com/museum](http://www.intel.com/museum)

Copyright © 2008 Intel Corporation. All rights reserved. Intel, Intel logo, Celeron, Intel386, Intel486, i386, i486, Intel Core, Intel Xeon, Itanium, and Pentium are trademarks of Intel Corporation in the U.S. and other countries. \*Other names and brands may be claimed as the property of others.

0108/TM/LAI/HP/20K 308301-002US